Contents lists available at ScienceDirect

Performance Evaluation

journal homepage: www.elsevier.com/locate/peva

Polynomial analysis algorithms for free choice Probabilistic Workflow Nets[☆],[☆☆]

Javier Esparza[a,*], Philipp Hoffmann[a], Ratul Saha[b,*]

[a] *Technische Universität München, Boltzmannstr. 3, 85748 Garching, Germany*
[b] *National University of Singapore, Singapore*

ARTICLE INFO

Article history:
Available online xxxx

Keywords:
Workflow Petri nets
Expected reward
Free-choice Petri nets
Confusion-free Petri nets

ABSTRACT

We introduce Probabilistic Workflow Nets (PWNs), a model extending confusion-free workflow Petri nets with probabilities. We give PWNs a semantics in terms of Markov Decision Processes (MDPs) and introduce a reward model. We show that the expected reward of a complete execution of a PWN is independent of the scheduler used to resolve the nondeterminism of the MDP, which allows one to choose a suitable scheduler for its computation. However, this feature does not lead to a polynomial algorithm, and in fact we prove that deciding whether the expected reward exceeds a given threshold is `PSPACE`-hard.

To alleviate this high computational cost, we extend previous work on property-preserving reductions of non-probabilistic workflow nets. We introduce reduction rules for PWNs, and prove that they preserve the expected reward. The rules allow us to simplify the workflow before constructing its MDP. We then consider the subclass of free-choice PWNs, whose non-probabilistic counterpart has been extensively studied. Using a previous result on the power of the rules for this class, published by us in FASE'16, we derive a polynomial-time algorithm in the size of the PWN for the computation of the expected reward. In contrast, algorithms based on constructing the MDP require exponential time. We report on a sample implementation of the reduction algorithm and on its performance on a collection of benchmarks.

Finally, we present two extensions of our work. First, we show that our reduction rules can also be used to compute the expected reward parametrically, that is, as a function of parameters related to the probabilities and rewards of the transitions. Second, we discuss the extension of PWNs to workflow nets that are not confusion-free, and show that some of our results still hold.

1. Introduction

Workflow Petri Nets are a class of Petri nets for the representation and analysis of business processes [1–3]. They are a popular formal back-end for different notations like BPMN (Business Process Modeling Notation), EPC (Event-driven Process Chain), or UML Activity Diagrams.

[☆] A previous version of this paper was presented at the 13th International Conference on Quantitative Evaluation of SysTems (QEST 2016).

[☆☆] This work was partially funded by the DFG Project 5090812 (Negotiations: Ein Modell für nebenläufige Systeme mit niedriger Komplexität).

[*] Corresponding authors.
E-mail addresses: esparza@in.tum.de (J. Esparza), hello@ratulsaha.com (R. Saha).

https://doi.org/10.1016/j.peva.2017.09.006
0166-5316/

Please cite this article in press as: J. Esparza, et al., Polynomial analysis algorithms for free choice Probabilistic Workflow Nets, Performance Evaluation (2017), https://doi.org/10.1016/j.peva.2017.09.006.

There is recent interest in extending these notations, in particular BPMN, with the concept of cost (see e.g. [4–6]). The final goal is the development of tool support for computing the worst-case or the average cost of a business process. A sound foundation for the latter requires to extend Petri nets with probabilities and rewards. Since Petri nets can express complex interplay between nondeterminism and concurrency, the extension is a nontrivial semantic problem which has been studied in detail (see e.g. [7–9] for untimed probabilistic extensions and [10,11] for timed extensions), and is not yet settled.

This paper is based on the following two observations:

- Many workflow models of business processes avoid a situation called *confusion*, in which concurrent transitions and transitions in conflict "interfere" [12]. For example, the Workflow Graph formalism, also called free-choice Workflow Nets [1,13–15], does not allow to describe workflows with confused markings, and 1386 of the 1958 nets in the most popular benchmark suite in the literature are free-choice Workflow Nets [16].
- Giving a semantics to probabilistic Petri nets is much simpler in the confusion-free case; in particular, different semantics collapse for this subclass [7,8].

This suggests to investigate analysis algorithms for probabilistic confusion-free workflow nets. Following ideas of [10,11] we give these nets a natural semantics in terms of Markov Decision Processes (MDP) with rewards. In a nutshell, at each reachable marking the enabled transitions are partitioned into disjoint *conflict sets*; transitions in the same conflict set are in conflict, while transitions in different conflict sets are concurrent. A scheduler selects one of the conflict sets, and then one of the transitions of this set is chosen probabilistically.

In our first contribution we prove that the expected reward of a confusion-free workflow net is independent of the scheduler resolving the nondeterministic choices, and so we can properly speak of *the* expected reward of the net. This leads to a simple algorithm for the computation of the MDP: Choose a memoryless scheduler, compute the Markov chain corresponding to it, and determine its expected reward by solving a system of linear equations. Since the size of the Markov chain can be exponentially larger than the size of the workflow net (a fact usually called the *state-explosion problem*), this algorithm has exponential worst-case complexity. We prove that no polynomial algorithm exists unless $\mathtt{P} = \mathtt{PSPACE}$, because the problem of determining whether the expected reward exceeds a given bound is $\mathtt{PSPACE}$-hard.

Our second contribution is a *reduction algorithm* that palliates the state-explosion problem. The algorithm repeatedly applies a set of *reduction rules* that simplify the workflow while preserving its expected reward. The rules are an extension to the probabilistic case of a set of rules for free-choice Colored Workflow Nets recently presented in [15]. They merge two alternative tasks, summarize two consecutive tasks into one, and replace a loop with a probabilistic guard and an exit by a single task.

In our third contribution we study the special case of free-choice workflow nets. While they are a proper subclass of the confusion-free workflow nets, they are also widely used. In particular, the evidence supplied above for the abundance of confusion-free workflow nets also applies to free-choice workflow nets. With the help of a result from [15], we give a polynomial-time algorithm to compute the expected reward. The algorithm either stops with the conclusion that the expected reward is infinite, or reduces the initial free-choice workflow net to a workflow net with only one transition; since the rules preserve the expected reward, the reward of this transition is equal to the expected reward of the initial workflow net. We report on a prototype implementation, and on experimental results on the benchmark suite of [16], which contains workflows derived from industrial business processes. We compare our algorithm with the different algorithms based on the construction of the MDP implemented in PRISM [17].

In the last sections of the paper we consider two extensions of our results. First, we show that our reduction algorithm can also be used to compute *parametric* expected rewards. In the parametric case the probabilities or rewards associated to transitions are not numbers, but expressions involving parameters, and the problem is to compute the expected reward as a function of them. This problem was first studied in [18] for Markov chains, and for other Markov models and rich specification formalisms in [19–22]. We show that our reduction rules for probabilistic workflow nets can be lifted to the parametric case.

In the final section we consider probabilistic workflow nets that are not necessarily confusion-free. We show that our semantics is mathematically valid, and discuss some of its properties. In particular, we give a simple example of a workflow net where the expected reward depends on the scheduler. We show that some of our results can still be transferred to the general case.

Structure of the paper

Section 2 recalls basic definitions about workflow Petri nets. Section 3 introduces Probabilistic Workflow Nets with Rewards (PWNs) as a probabilistic extension of confusion-free workflow nets, and gives them a semantics in terms of MDPs. Section 4 shows that PWNs have the same expected reward under all schedulers, deduces an algorithm for computing the expecting reward, and proves that no polynomial algorithm exists unless $\mathtt{P} = \mathtt{PSPACE}$. Section 5 presents our reduction rules and proves that they preserve the expected reward. Section 6 introduces free-choice PWNs and obtains a polynomial time reduction algorithm. Section 7 reports on our implementation of the reduction procedure and our experiments. Section 8 introduces parametric PWNs and lifts our reduction rules to this case. Section 9 discusses probabilistic workflow nets that are not necessarily confusion-free. Finally, Section 10 presents some conclusions.

Please cite this article in press as: J. Esparza, et al., Polynomial analysis algorithms for free choice Probabilistic Workflow Nets, Performance Evaluation (2017), https://doi.org/10.1016/j.peva.2017.09.006.